Salinity Monitoring Using Remotely Sensed and Other Spatial Data

Furby, S.L., Caccetta, P.A. and Wallace, J. F.
CSIRO Division of Mathematical and Information Sciences

Abstract

Australia's south west agricultural region is severely affected by dryland salinity, which has damaging effects on agricultural productivity, built infrastructure and conservation values. Salinisation of land and waterways is the highest priority environmental issue in the region, but accurate information on the extent of the problem over the region can not be provided by conventional means.

Spatially explicit maps of salinity, its change over approximately 10 years, and of areas at risk from salinity have now been produced for the entire region (24 million hectares) using time series Landsat imagery and accurate DEM data. This has been achieved in the 'Land Monitor' project, which operationalised methods developed in earlier research projects. Ground data provided by experts was used to guide the analyses and for validation.

The key to providing the mapping and monitoring information was the application of enhanced maximum likelihood classification to a sequence of images over several years. The classifications were then combined in appropriate multi-temporal models to reduce errors and to produce acceptable change maps. This paper describes the mapping and monitoring methodology.

Keywords: Remote Sensing, Salinity, Monitoring, Multi-temporal, Landsat.

1. Introduction

The Western Australian State of Environment Report (1998) identified land salinisation, salinisation of inland waters and maintenance of biodiversity as the three highest priority environmental issues in Western Australia. Since settlement by Europeans, the south-west of Western Australia has been extensively cleared for agriculture. Ground waters in the region are generally saline, and the replacement of deep-rooted native perennial vegetation with shallow-rooted annual crops and pastures has altered the hydrological balance. Ground waters are generally rising with consequent increases in dryland and stream salinity. Productive land has been lost and water supplies to rural and urban communities have been affected. Salinity also causes damage to rural infrastructure including buildings, roads and rail links, and services such as underground sewerage, telephone and gas and electricity supply systems. Environmental impacts on remnant vegetation and former fresh water systems have been severe with consequent impacts on biodiversity values (Halse et al, 2003 and references therein).

Prior to the work described here, the state of knowledge of the extent and changes in dryland salinity was poor (Ferdowsian et al, 1996). It was estimated that about 1.8 million hectares in Western Australia were already salt-affected and that this area could double in the next 15 to 25 years and then double again before reaching equilibrium (reference). Historically, two methods have been used to obtain information on salinity: surveys requesting land managers provide estimates of land use and condition, and manual interpretation of aerial photography (George, 1990). Questionnaire responses from farmers on salinity were recognised as incomplete, inaccurate and non-spatial. Interpretation of aerial photography had been carried out in some areas, but the results were recognised as subjective, labour-intensive and incapable of broad application over the region. In addition, aerial photography is acquired irregularly over the region and generally collected in the dry summer months when only the most severely affected areas are visible.

Classification of single-date Landsat imagery had also been evaluated and research conducted which identified the spring growing season as the optimal time of year for discrimination of salinity (Furby et al, 1995). Salinity affects growth of crops and pastures and the health of native perennial vegetation. Severely-affected areas appear as bare scalds or areas of dead trees. In marginally affected areas, pastures and crops evidence reduced productivity and may be replaced by salt tolerant species such as barley grasses (reference). Using single-date growing-season imagery alone, severely affected areas were mapped well, with some errors of commission from areas which were bare from other causes (Furby et al., 1995). Attempts to classify marginally-affected areas resulted in large errors of commission. Salinity is only one possible cause of low productivity; in particular, farmer's management decisions applied in each season can result in similar effects for non-saline areas. As a consequence of these errors, single-date classifications were of limited value, and it was recognised that simple approaches to monitoring and change detection (such as post-classification differencing) would not be useful. However, the farming system is based on crop/pasture rotation. Over successive seasons, the temporal pattern of these 'management effect errors' is quite different from salinity effects on productivity. Salinity persists through time and may become more severe while management effects tend to be transient.

Based on this experience and knowledge of the management system, an approach to salinity mapping and monitoring using a sequence of Landsat Thematic Mapper (TM) images together with digital elevation models (DEMs) and ground-truth data was developed. A key to the success of the approach was the retention of measures of confidence from the classifications of individual images, and the joint processing of this information over the time series in an appropriate model. Information derived from the DEMs was also incorporated in the model. The result was improved accuracy of mapping which provided acceptable monitoring of change. Ground data provided by experts was used to guide the analyses and for validation. Using a different approach, spatially explicit maps of salinity hazard (i.e. areas at risk of salinity in future) were also produced using DEM derivatives and land cover variables derived from Landsat.

The methodology has been implemented operationally in the southwest agricultural area of Western Australia (a region of 24 million hectares of land) as part of the Land Monitor project (Caccetta et al, 2000 and Land Monitor, 2008). The products are a baseline map of salt-affected land, a map of change since the baseline interval and a map of areas potentially at risk of salinity in the future. The results are widely distributed to government agencies (e.g the Western Australian Departments of Agriculture and Food, environment and Conservation, Water and Planning and Infrastructure and the WA Land Information Authority) and land managers.

The following section of this paper describes the key elements of the salinity mapping and monitoring methodology. The third section presents some typical results from the Land Monitor project. The salinity hazard mapping is not discussed here (see Land Monitor, 2008).

2. Salinity Mapping and Monitoring Methodology

The steps in the salinity mapping and monitoring process were:

1. Acquire relevant image and supporting spatial data
2. Ortho-rectify and calibrate the image data
3. Extract relevant terrain variables from the DEM data
4. Stratify the region into zones with little regional variation
5. Single–date land use/condition classifications
6. Multi-temporal classification.

The key aspects of each of these stages are discussed below.

2.1 The data

Given the area to be covered, the required spatial resolution, and the monitoring objectives, Landsat TM satellite imagery was the only suitable choice to provide data relating land condition. Although the rate of spread of salinity depends on a number of factors, it was necessary to be able to monitor changes of the order of 25 – 100m over a 3-5 year interval. Australia has received and archived Landsat TM imagery routinely since 1988. Approximately 18 Landsat scenes are required to cover the south-west agricultural area in Western Australia.

As noted above, images from the time of maximum green vegetation cover (spring) are optimal for mapping salt-affected land (Wheaton et al, 1992). This time of year provides maximum discrimination between good condition land (productive pasture or crop growth) and land that is potentially salt-affected. Other research has shown that a sequence of two to three images in successive growing seasons (years) are required to discriminate between persistent land condition problems, such as salinity, and ephemeral conditions, such as management and seasonal variation (Furby et al. 1995). Hence, a series of three images is required to produce a reasonable baseline for monitoring. In practice and to provide the monitoring interval, at least six images were acquired for each area from 1988 to the time of processing.

Landform is a significant factor in groundwater movement and the expression of salinity. In general, areas on hilltops and ridges are less likely to be saline compared to low lying areas such as depressions and valley floors (Caccetta, 1997). This knowledge can be used in the analysis provided that consistent data on relative landform position can be created from the DEM (2.3 below).

Ground-truth data was required to train the classifications and to provide an assessment of the accuracy of the resulting maps. Regional experts supplied this information, and later performed the accuracy assessments. The training data were provided in the form of areas marked on aerial photographs, farm plans, maps or images and included the extent of salt-affected areas and a description of the ground cover types within the affected areas.

2.2 Ortho-rectification and calibration

Change detection requires that the images be well registered to each other. An ortho-rectified base mosaic was created for the region using images from a single time period using ground control points. The remaining images in the sequence were registered to the base mosaic. Cross-correlation feature matching techniques were used to increase the speed and accuracy of the co-registration of the images. Ideally, all images would be calibrated to standard reflectance units. However, when comparing images to detect change, it is sufficient to convert the raw digital counts to be consistent with a chosen reference image. Such a 'like values' calibration procedure using invariant targets was implemented (Furby and Campbell, 2001).

2.3 DEM-derived variables

The existing contour data over the region was inadequate for the purpose and so a new DEM was produced using soft-copy automated photogrammetry techniques at a cost of approximately AUD$3m (Land Monitor, 2008). This Land Monitor DEM has a spatial resolution of 10m and a height accuracy of +/- 1m. Terrain variables were derived from the DEM (Caccetta et al, 2000) using multiple flow algorithms (Quinn et al, 1991). Pre-processing included pit filling using the methods of Gratin and Soille (1993). The derived variable upslope area was partitioned to provide a consistent surrogate landform units (hilltops, ridges and upper slopes, upper valleys, lower valleys, and broad valleys). The landform partitioning provides strong prior evidence of which parts of the terrain are likely to be/become salt-affected and which parts are not (Caccetta, 1997).

2.4 Stratification

Each Landsat TM satellite image covers an area approximately 200km by 200km. Within that area, variations in rainfall, farming practices, soils and geology affect the type and vigour of vegetation cover and the amount and type of salt-affected land. In turn, these factors affect the spectral signals of healthy and salt-affected land and the error rates expected in each zone. For this reason the image data were stratified into 'uniform' zones for classification based on visual assessment of the imagery and supporting ancillary data such as rainfall and soil type mapping. Typically the classification results obtained for one zone did not extrapolate to other zones. The results for each stratification zone were mosaiced at the end of the processing.

2.5 Single-date land use / condition classification

Within each image and zone, supervised maximum likelihood classification (MLC, Rao, 1966) was applied to map saline and non-saline land cover types. At least two saline classes were identified: 'severe' (mainly bare or halophytic vegetation), and 'marginal' which included areas of reduced productivity and barley grasses. The enhanced MLC approach described in Wallace and Campbell 1989 was used. Class definitions and signatures were based on spectral separability using canonical variate analysis (CVA) (Campbell and Atchley 1981). In addition, the allocation procedures calculate and retain the posterior probabilities (relative likelihood) for each of the spectral information classes. In a typical MLC analysis, each pixel is assigned to the class with the highest posterior probability. Here the posterior probabilities from each image date give a measure of confidence; these measures over the time sequence are combined with landform information to refine the classification.

First pass results for each zone and date were reviewed. Where errors or large areas of uncertainty were found, the analysis and classification were repeated, often with additional training sites, until no further significant improvement could be obtained.

2.6 Multi-temporal classification

The accuracy of single-date classification varied with the land cover class. Severe salinity was typically well mapped with some errors of commission, while marginal salinity was less accurately mapped with large errors of commission. Classification accuracies also varied with season and time of image acquisition. If the rainfall season was poor or an image was not available at the optimal time due to cloud, then poor vegetation responses were widespread and accuracy was reduced. In most cases, the single-date classifications were not accurate enough to use as a resource management tool.

Knowledge of classification accuracies and of temporal dependencies between classes was used to improve the results. Results from non-optimal dates were also downweighted. Poor cover that persists for only one or two growing seasons was considered more likely to be a result of management or seasonal affects. Only if poor cover persists through multiple growing seasons, was the area considered likely to be saline. Position in the landscape, derived from the DEM, was also used as a prior to modify the probability of such areas being labelled as saline.

The conditional probability network (CPN) is a Bayesian network, which provides a computational framework to combine uncertain classified satellite image data from several growing seasons with the landform data. CPNs allow for the assessment and propagation of uncertainty from multiple sources of data of varying quality or accuracy (Caccetta, 1997). The scheme for combining data was based upon techniques presented in Lauritzen and Spiegelhalter (1988) and Lauritzen (1992).

Useful properties of the Bayesian network approach for the application included:

- propagation of uncertainties in inputs and calculation of uncertainties in outputs;
- production of hard (class label) and soft (probability) maps;
- estimates for missing data (e.g cloud-covered image data in a particular year) are calculated using all available information; and
- well-developed statistical tools for parameter estimation exist.

The network can be represented as a graph, where the nodes of the graph represent random variables and the edges of the graph represent conditional dependency relationships between the variables. An example of a network for mapping and monitoring salinity is shown in Figure 1. The circles and rectangles represent the nodes of the graph. The rectangles represent variables that are observed and the circles represent variables that are not directly observed, in this case the true salinity map at each date. The top row of rectangles represents the estimated salinity map from the classification of the images for each growing season. The probability of belonging to each cover class, rather than a hard class label, is input to the network to reflect the relative certainty of the classification. The bottom rectangle represents the landform unit map derived from the DEM, which is entered as a hard class label. Values for the unobserved variables can be inferred from the other variables.

The graph edges or 'arrows' represent probabilistic relationships (conditional dependencies) between the variables. Observing a particular value for a variable provides some information about all the other variables to which it is connected (i.e. knowledge of landform class gives some information about likelihood of salinity). The rules, or relationships, between the variables are expressed in terms of conditional probability tables.

- The vertical arrows in Figure 1 represent two-dimensional tables which describe the accuracy (uncertainty) of the single-year classification maps; entries record Pr(true class|observed class).
  For example, in a dry year we would expect
    Pr (Productive | classified productive) is high (0.95), while
    Pr (Salinity | classified marginal salinity) is relatively low (0.50)
  With typically five land cover classes, each vertical arrow represents a 5 by 5 table of estimated accuracies.
- Temporal rules (horizontal arrows in Figure 1) link the true land condition state through time; i.e. Pr(state at year Y+1 | state at year Y). For example, we know that transitions from saline to

productive (recovery) are truly rare. However, there is a low probability of transition of productive land to saline in any one period and that probability depends on landform position. So, in a valley, for example, experts may estimate

Pr (Productive at Y+1 | Productive Y) = 0.97, while

Pr (Saline at Y+1 | Productive Y) = 0.03.

On an upper slope, transition to salinity would be much less likely. The dependency on landform (diagonal arrows in Figure 1) effectively means that 3-dimensional tables are required to parameterise these links in the network.

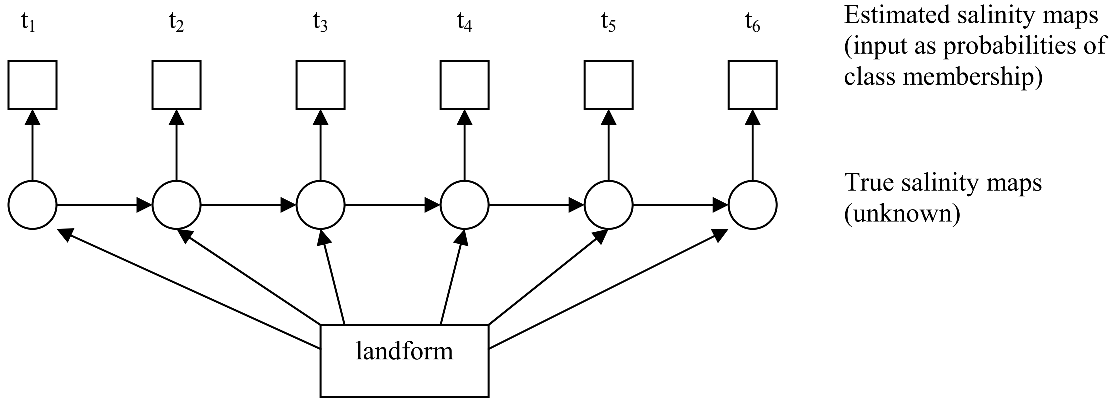

Figure 1: An example of a conditional probability network for salinity mapping and monitoring. In this example, a sequence of six classifications have been produced from six Landsat TM images and are combined with landform information derived from a digital elevation model.

In the network illustrated (Figure 1), the input data to the CPN for each pixel are its relative landform position and the posterior probabilities of land cover class at all dates. The outputs are a set of posterior probabilities at all dates, which have been modified by the relationships in the network. It is important that the network represents a realistic set of relationships between variables and that plausible values for the tables can be estimated. Here the classification at each year does provide estimates of the true land cover state at that time with estimates of accuracy in the 'vertical arrow' tables. The temporal dependencies of land condition state, modified by landform, also represent understanding of the system. These dependencies may vary with geographic zone. Therefore, the network above is parameterised zone by zone. A more complex network including a 'zone' node and its dependencies could be set up to process a number of zones simultaneously.

It is necessary to estimate reasonable conditional probability values to parameterise the network tables. Here estimates for the classification accuracy tables (vertical arrows in Figure 1) were provided from ground data, which were not used in the initial analysis. Initial values for temporal and landform rules were estimated using an EM (expectation – maximisation) algorithm (Dempster et al, 1977; Lauritzen 1995). These values were sometimes modified to better reflect expert knowledge, e.g. to reduce the probability of salinity on hilltops.

Figure 2 illustrates the effect of the multi-temporal processing on the classification results. The poor season image showed large areas over-classified as 'marginally saline' based only on the single-date data. Much of the land classified marginal in the 'poor' season was reclassified to productive based on the whole temporal sequence and landform information. Even the allocation of some bare areas with high probability of severe salinity in the above average season was changed based on the full sequence and landform information. Differencing the two refined classifications produced far less commission errors and hence formed a monitoring product that could be used to inform resource management decisions.

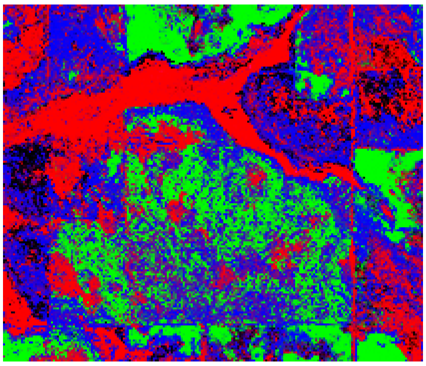

Single-date classification – poor season

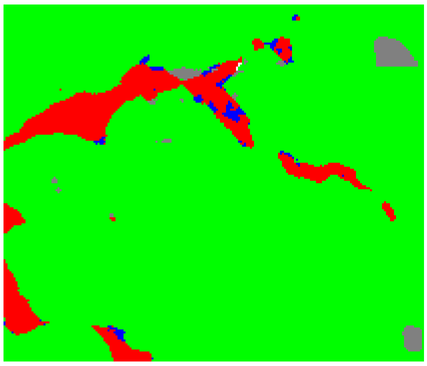

Multi-temporal classification – poor season

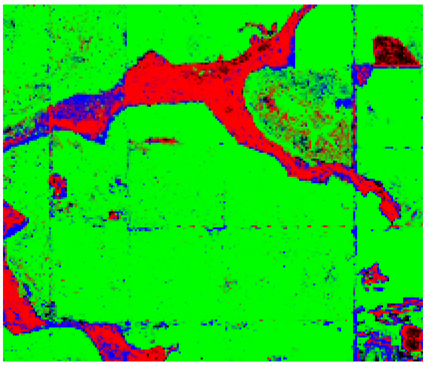

Single-date classification – above average season

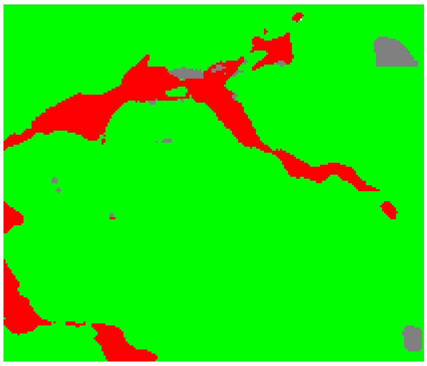

Multi-temporal classification – above average season

Figure 2: Single-date classifications for poor and above average seasons (left) and corresponding multi-temporal classifications (right). In the single-date classifications probabilities of productive cover are shown in green, probabilities of marginal cover are shown in blue and probabilities of severe salinity are shown in red. In the multi-temporal classifications class labels rather than probabilities are shown in the same colours. Grey corresponds to areas mapped as remnant perennial vegetation (bush).

3. Results from the Land Monitor Project

The salinity mapping and monitoring methodology has been routinely applied to the south-west agricultural area in Western Australia as part of the Land Monitor project. The accuracy of the salinity extent and maps was determined by independent ground assessment of regular grid points on sample areas from each zone. The full results are described in a series of reports (Land Monitor, 2008) from which an example from a typical zone is reproduced below. These particular tables are reproduced from McFarlane et al, 2004.

The tables show summaries of the salinity mapping accuracy as site counts (Table 1), "user's accuracy" (Table 2) and "producer's accuracy" (Table 3). The overall accuracy was 95% with errors of omission and commission being present. For each region, a description of the major sources of error were noted. Generally, errors of commission included sites such as degraded bush in low lying areas, eroded catchments and new dams, which were not included in the dam mask. Errors of omission included underestimates of the extent of hillside seeps, and saline sites (or edges of saline sites), which supported productive plant growth. The size of the surface expression of the salinity relative to the image pixel size was also a factor in the success of the process. If the extent of the salt-affected area was less than an image pixel (25m for Landsat TM imagery) then it was unlikely to reduce the vegetation cover sufficiently to be mapped as potentially saline. For example, narrow areas of emerging salinity were noted as one source of omission errors in the processing.

Table 1: Site counts (pixels) obtained from field validation (example only)

		Image Label		
		Non saline	Saline	Total

Ground Label	Non saline	98	1	99
	Saline	6	26	32
	Total	104	27	131

*Table 3: Percentage of ground classes mapped (*errors of omission, ** errors of commission)*

		Image	Label	
		Non saline	Saline	Total (%)
Ground Label	Non saline	99	1**	100
	Saline	19*	81	100

*Table 2: Percent accuracy of mapped classes (*errors of omission, ** errors of commission)*

		Image	Label
		Non saline	Saline
Ground Label	Non saline	94	4**
	Saline	6*	96
	Total (%)	100	100

5. Conclusions

Using the methods above, new and accurate information on extent and change of saline land has been produced over a large area. The results showed how integration of information from multiple image dates and DEMs can be combined to improve the classification accuracy. Keys to this development were the understanding of the physical processes and of dependencies between the datasets, recognition of error in individual classifications, and handling of this uncertainty in an appropriate mathematical framework. Inspection of interim results by land management experts, iterative refinement and validation have been important steps in producing optimal results at each stage. These results provided confidence in the accuracy and the use of the products. Validation produced accuracy statistics in all areas, and identified and documented certain types of errors and underestimates in particular areas.

The Land Monitor method has estimated that the area currently affected by salinity in Western Australia is about 1 million hectares (in 1996) and the annual rate of increase is about 14,000 ha. In addition to regional summaries, spatially explicit maps of salt-affected land are now available that directly inform catchment management decisions. Further refinements and updates of Land Monitor methodology and area estimates are warranted as it is eight years since the last estimate and saline areas may have increased by a further 100,000 ha, or 0.3%, in this time.

Acknowledgements

The mapping in Western Australia was performed as part of the Land Monitor project – a multi-agency state government initiative (funded in part by the National Heritage Trust) composed of Agriculture and Food Western Australia, the Western Australian Land Information Authority (Landgate), Main Roads Western Australia, the Department of Environment and Conservation, the Department of Water and the Commonwealth Scientific and Industrial Research Organisation. The results presented were produced with considerable ground-truthing effort from experts in the member agencies.

References

Caccetta P. A. (1997), Remote sensing, geographic information systems (GIS) and Bayesian knowledge-based methods for monitoring land condition, PhD Thesis, *School of Computing, Curtin University of Technology*.

Caccetta, P.A, Allen, A. and Watson, I. (2000). The Land Monitor project. *Proceedings of the 10th Australasian Remote Sensing Conference*, Adelaide, Australia, 2000.

Campbell, N.A. and Atchley, W. R. (1981). The geometry of canonical variate analysis. Syst. Zool., 30(3), 268-280.

Dempster, A. P., Laird, N. and Rubin, D. B. (1977), ‘Maximum likelihood from incomplete data via the EM algorithm’, *Journal of the Royal Statistical Society*, B, Vol 39, pp. 1-38.

Environment Western Australia 1998: *State of the Environment Report,* Published by the Department of Environmental Protection, Government of Western Australia, July, 1998.

Ferdowsian, R., George, R., Lewis, F., McFarlane, D., Short, R. and Speed, R (1996), The extent of dryland salinity in Western Australia. *Proceedings of the 4th National Conference and Workshop on the Productive Use and Rehabilitation of Saline Lands*, Albany, WA, March 1996, pp. 89-97.

Furby, S. L. and Campbell, N. A. (2001). Calibrating images from different dates to 'like value' digital counts. *Remote Sensing of Environment*,77, p186-196.

Furby, S. L., Wallace, J. F., Caccetta, P. A. and Wheaton, G. A. (1995). Detecting and monitoring salt-affected land. *Report to LWRRDC, Project CDM1*.

George, R (1990). The 1989 saltland survey. *J. Agric. West. Aust.*, 31, 159-166.

Gratin, C. and Soille. P. (1993), Short Communication: An Efficient Algorithm for Drainage Network Extraction on DEMs, Centre de Morphologie Mathematique, Ecole Nationale Superieure des Mines de Paris, 35, rue Saint-Honore F-773005 Fontainebleau Cedex, France.

Halse, S.A., Ruprecht, J.K., and Pinder, A.M. 2003. Salinization and prospects for biodiversity in rivers and wetlands of south-west Western Australia. *Australian Journal of Botany* 51, 673-688.

Land Monitor (2008), Land Monitor [Internet documents]. *Land Monitor*, Perth [cited January 2008]. Available from http://www.landmonitor.wa.gov.au

Lauritzen, S.L. and Spiegelhalter, D.J. (1988). Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society*, B, Vol 50, No 2, pp 157-224.

Lauritzen, S.L. (1992). Propagation of probabilities, means and variances in mixed graphical association models, *Journal of the American Statistical Association*, Vol. 87, No. 420, pp. 1098-1108.

Lauritzen, S.L. (1995). The EM algorithm for graphical association models with missing data, *Computational Statistics and Data Analysis,* 19, pp. 191-201.

McFarlane, D. J., George, R. J. and Caccetta, P. A. (2004), The Extent and Potential Area of Salt-affected Land in Western Australia Estimated Using Remote Sensing and Digital Terrain Models, *Engineering Salinity Solutions*, 9-12 November, Perth, Western Australia.

Quinn, P., Beven, K., Chevallier, P. and Planchon, O. (1991), The prediction of hillslope flow paths for distributed hydrological modelling using digital terrain models, *Hydrological Processes*, Vol. 5, No. 1, pp. 59-79.

Rao, C. R. (1966). *Linear statistical Inference and its Applications*, John Wiley and Sons Inc, New York, Second edition, pp 289-191.

Wallace, J. F. and Campbell, N. A. (1989), Analysis of remotely sensed data*. In Remote Sensing of Biosphere Functioning* (Eds, R J Hobbs and H A Mooney) 1990. Springer-Verlag, New York.

Wheaton, G. A., Wallace, J. F., McFarlane, D. J. and Campbell, N. A. (1992), Mapping salt-affected land in Western Australia, *Proceedings of the 6th Australasian Remote Sensing Conference*, vol 2, pp. 369-377.